INCORPORATION BY REFERENCE

Pages 35-308 comprises the following Appendices which are hereby incorporated herein by reference in their entirety:

APPENDIX A

OWL NETWORK ARCHITECTURE

Overview.

Norand's open wireless LAN (OWL) architecture is designed to facilitate wireless communications at the MAC sub layer of the ISO protocol stack. An OWL radio network can function as a stand-alone LAN or it can function as a subnet in an 802 LAN to provide wireless access to wired 802 subnets. An 802 LAN may include multiple wired 802 subnets and OWL subnets. Figure 1 shows an example 802 LAN which includes an OWL subnet. The OWL subnet (i.e. subnet 4) includes the OWL radio network (i.e. subnet 2) and a "secondary" 802.3 subnet (i.e. subnet 3).

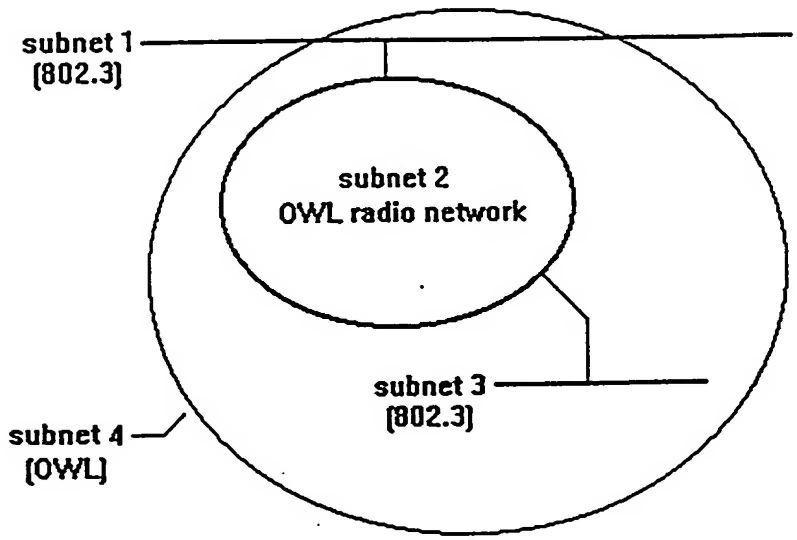

figure 1.

Figure 2 shows an example 802 LAN, similar to the LAN in figure 1, with an expanded view of the OWL radio network. Subnet 1 is not part of the OWL subnet, however it provides a distribution LAN for the OWL subnet. An OWL radio network provides wireless access to the 802 LAN for mobile radio-equipped computers (MRCs). An OWL radio network can also provide a wireless transparent bridge between wired 802 subnets (i.e. an OWL subnet can include a wired 802 subnet). Any node in an 802 LAN, which includes an OWL subnet, can communicate with any other node, at the logical link control (LLC) sub layer of the data link layer. In figure 2, remote station 1 can communicate with either MRC or remote station 9. MRC 6 can communicate with MRC 8 or either remote station.

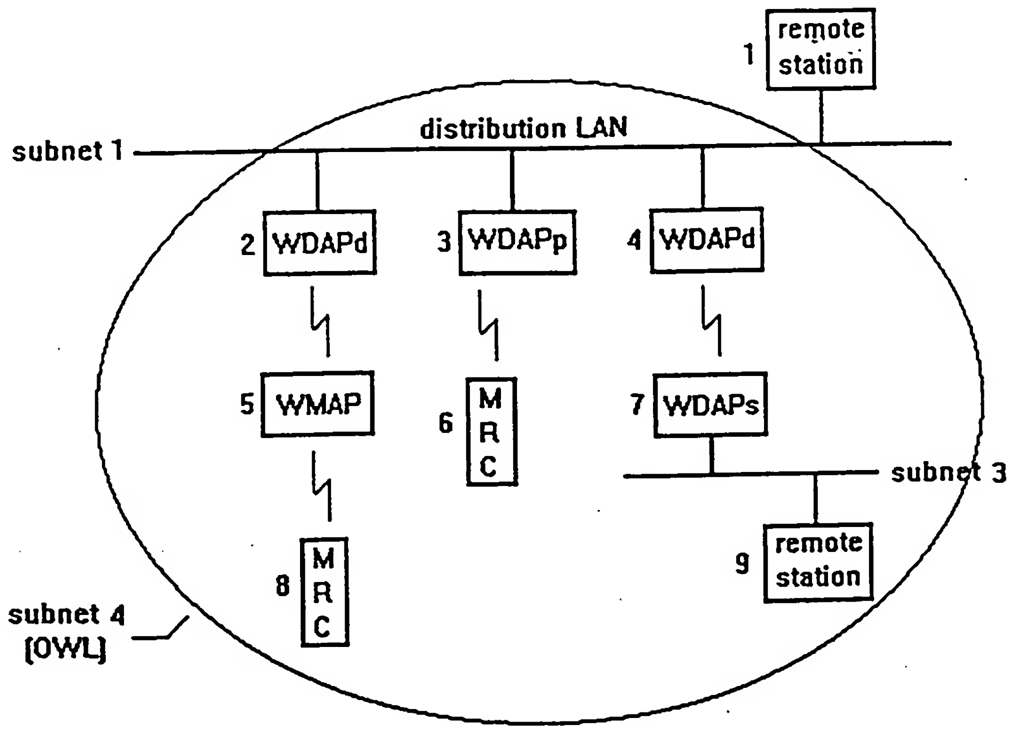

Figure 2.

The IEEE 802.11 committee has defined two basic types of wireless networks - hierarchical and ad hoc. Hierarchical networks contain radio-equipped access points which provide a centralized control function for a given radio coverage area. All communications pass through the access point. The access point also provides access to a wired LAN. A hierarchical network may contain multiple access points which provide an extended seamless radio coverage area. Mobile computers can roam from one access point coverage area to another. Ad hoc networks facilitate peer-to-peer communications in the absence of a central control point. This document is primarily directed toward hierarchical networks.

The OWL protocol stack is contained in the MAC sub layer of the ISO protocol stack. An OWL MAC (i.e. in a terminal node) provides MAC sub layer services to the LLC sub layer of the ISO data link layer. The OWL MAC is subdivided into 4 sub layers: MAC-D, MAC-R, MAC-Q, and MAC-S.

MAC-D - The MAC-D sub layer is analogous to the data link layer in the ISO protocol stack. The MAC-D layer provides data link services to the MAC-R layer. It is responsible for channel access control and the reliable transmission of MAC-R PDUs across a single link in the OWL network. The MAC-D sub layer is specific to the link type (i.e. radio, ethernet, etc.).

MAC-R - The MAC-R sub layer is analogous to the network layer in the ISO protocol stack. The MAC-R layer provides routing services to the MAC-Q layer. It is responsible for correctly routing MAC-R PDUs through the OWL subnet, which may include multiple hops and circular physical paths.

MAC-Q - The (optional) MAC-Q sub layer adds reliability to the radio network by retransmitting lost PDUs. The MAC-Q layer is responsible for discarding out-of-sequence and duplicate PDUs. The MAC-Q sub layer can be implemented as an entity in the MAC-R sub layer. MAC-Q entites exist at entry points to the radio network.

MAC-S - The (optional) MAC-S sub layer is responsible for providing services for security, compression, etc. MAC-S entities exist at entry points to the OWL radio network.

A logical OWL node is a MAC-R addressable entity in an OWL radio network. An OWL node can be one of two types: 1) a terminal node or 2) a relay node. Terminal nodes are end points in the network; relay nodes forward PDUs at the MAC-R sub layer. Figure 3 shows MAC protocol stacks for both node types. The arrows represent the flow of data between MAC sub layers in each node type. (The upper layers in the relay stack are used to process PDUs addressed to the relay node.)

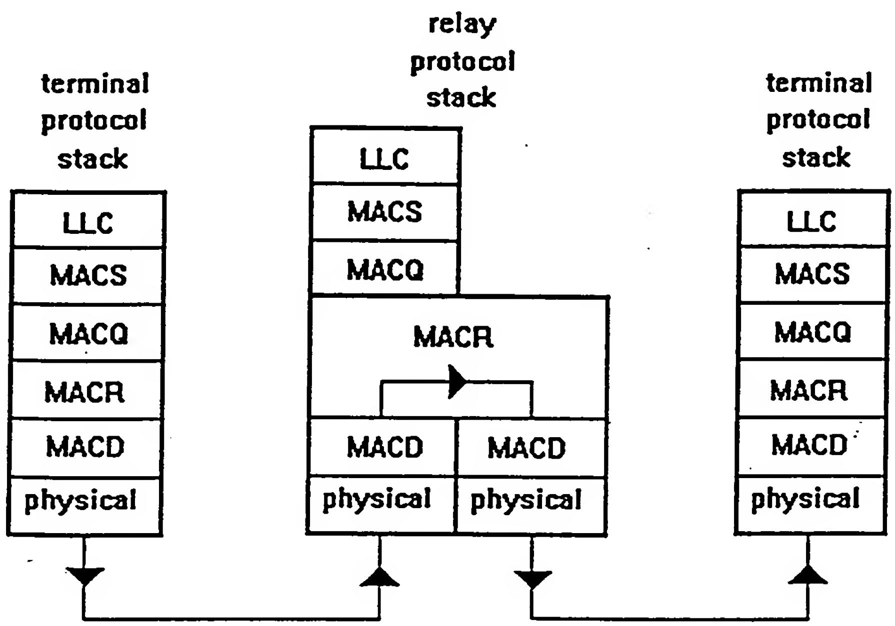

figure 3.

A wireless domain access point (WDAP) is an OWL bridge which is used to bridge a radio subnet to a wired 802 subnet. A WDAP contains a bridge protocol stack. Figure 4 shows the MAC protocol stack for a WDAP. Note that the bridge protocol stack contains a relay protocol stack. The 802.3 MAC-D sub layer is used to send OWL PDUs over an 802.3 link that is part of the OWL radio network. The MAC-Q and MAC-S sub layers serve as proxy MAC-Q and MAC-S entities for stations on the 802.3 sub net. The MAC-Q and MAC-S sub layers also service PDUs for the local WDAP 802 address.

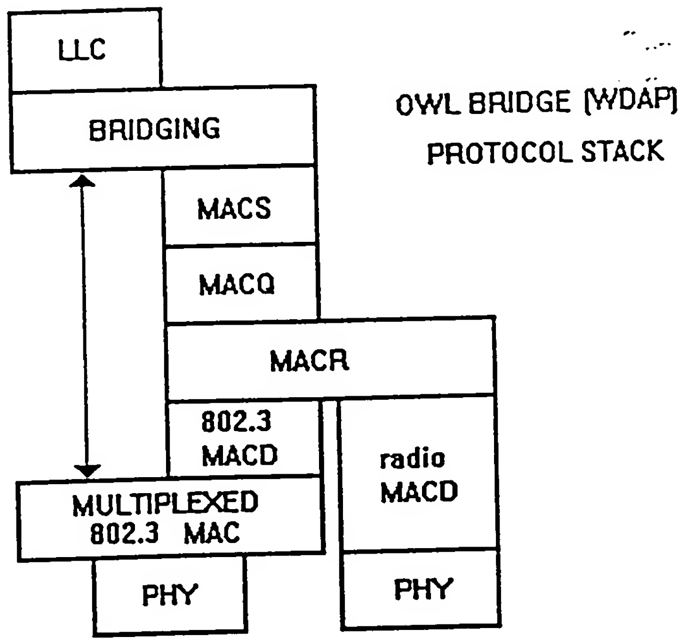

figure 4.

Figure 5 illustrates how data flows through a bridge protocol stack. The dotted line represents the path a PDU takes as it travels from a station on an 802.3 LAN to terminal 2 in an OWL radio network. The WDAP "bridges" the PDU from the 802.3 subnet to the radio subnet. The solid line represents the path a PDU takes as it travels from terminal 1 in the radio network to terminal 2 in the radio network. Since the path is contained in the radio network, the PDU does not have to be bridged.

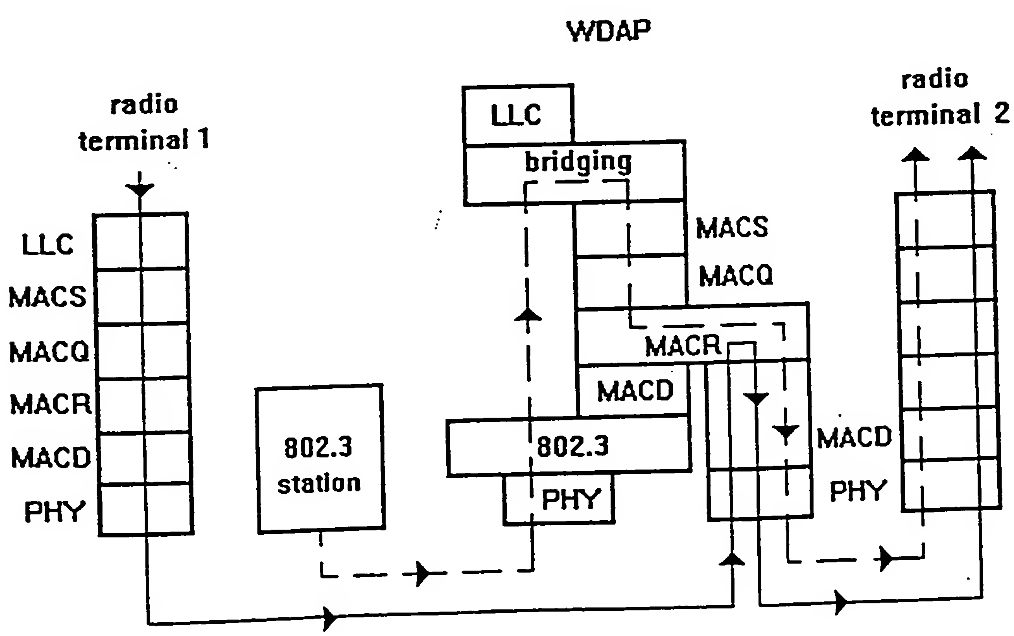

figure 5.

In general, PDUs are **bridged** across subnet boundaries; PDUs are **routed** within the radio network. A bridging entity in a WDAP uses a **forwarding database** to determine if a PDU should be bridged from one subnet to another subnet. A forwarding database contains a list of 802 address associated with each subnet to which the WDAP is attached. A MAC-R entity uses a **routing table** to determine how a PDU should be routed within an OWL subnet.

Network components and definitions.

802 LAN - a (possibly bridged) local area network which conforms to the IEEE 802 standards. For the purpose of this discussion, it is assumed that "802 LAN" refers to a LAN which contains wired 802.3 (ethernet) subnets and 1 or more OWL subnets.

802 subnet - a subnet in an 802 LAN which is not an OWL subnet.

OWL subnet - a subnet in an 802 LAN which includes an OWL radio network and 0 or more 802 subnets.

OWL Radio Network - An OWL subnet minus its wired subnets (see figure 1). An OWL radio network may include wired (i.e. 802.3) communications links. The OWL radio network consists of MAC-R addressable nodes and communications paths.

Mobile Radio-equipped Computer (MRC) - A mobile radio-equipped computer which contains an OWL terminal node.

Wireless Media Access Point (WMAP) - a radio-equipped base station which allows physical access to a wireless link in an OWL LAN. A WMAP may be connected to the radio network through a wired link or a radio link. A typical OWL radio network has multiple WMAPs with overlapping coverage areas. MRCs can roam between coverage areas. Except for possible timing issues, roaming has no effect on protocol layers above the MAC sub layer.

Wireless Domain Access Point (WDAP) - a logical access point to an OWL radio network. There are several types of WDAPs which are defined below. A WDAP is typically contained in a WMAP which is directly connected to a wired 802 subnet. The WDAP provides a bridge between the radio network and the wired subnet. A WDAP has a MAC-S and MAC-Q sub layer since it provides an entry point to the radio network. At any given time, one, and only one, WDAP provides access to a distribution LAN for a node in the OWL subnet.

OWL Node - A MAC-R addressable entity in an OWL radio network.

OWL Terminal Node - A MAC-R addressable OWL node which is an end point in an OWL radio network. A terminal OWL node is simply referred to as a **terminal** when the meaning is not ambiguous. A terminal has a MAC-S and MAC-Q sub layer since it provides an entry point to the radio network.

OWL Relay Node - A MAC-R addressable OWL node which is an interior node in an OWL radio network. MAC-R frames are routed through OWL relay nodes.

OWL Spanning Tree - An OWL spanning tree consists of a single root node, OWL relay nodes, terminal nodes, and edges, where a single edge logically connects two nodes for routing purposes. A **branch** is a logical path which contains 1 or more edges and the associated nodes. <u>MAC-R frames are routed along branches of a spanning tree</u>.

OWL Network Spanning Tree - All nodes in a hierarchical OWL subnet are organized into a network spanning tree for control purposes. A single network spanning tree constitutes an **OWL domain**. The root

of the network spanning tree contains a **primary WDAP**. Note that an 802 LAN may be attached to multiple OWL network spanning trees (domains).

OWL Access Spanning Tree - An access spanning tree is a sub tree in a network spanning tree. The root of an access spanning contains a **distributed or primary WDAP** and provides direct access to a **distribution LAN.**

Net ID - The Net ID identifies the set of nodes which belong to a single OWL domain - a network spanning tree or an instance of an ad hoc network. A hierarchical bit specifies whether the Net ID is for a hierarchical network or an ad hoc network. All nodes in an OWL domain share a common Net ID.

Super Root - the root of a network spanning tree. Multiple access points, attached to a distribution LAN, can negotiate to determine which node should function as the super root of a network. The super root is the node with the highest **super root priority**. The super root must have direct access to a distribution LAN. The super root is the **primary WDAP.**

Access Root - the root of an access spanning tree. An access root is a primary or distributed WDAP.

Distribution LAN - An 802 LAN segment which connects a wired subnet to the OWL subnet through the primary WDAP and 0 or more distributed WDAPs.

Distributed Root - the set of nodes which consists of the super root and all access roots. For a single OWL node, the distributed root can be viewed as the super root and the distributed WDAP which is providing access for the node to the distribution LAN.

Primary WDAP (WDAPp) - A single primary WDAP serves as the super root and provides a single control point for an OWL subnet. The primary WDAP has direct access to the distribution LAN. The primary WDAP "bridges" 802 frames from the distribution LAN to the OWL subnet and from the OWL subnet to the distribution LAN.

Distributed WDAP (WDAPd) - A distributed WDAP provides direct physical access to the distribution LAN. Distributed WDAPs exist within the domain of the primary WDAP. A distributed WDAP "bridges" 802 frames from the distribution LAN to the OWL subnet and from the OWL subnet to the distribution LAN.

Secondary WDAP (WDAPs) - An OWL subnet may include **remote 802 subnets** other than the distribution LAN. A single secondary WDAP serves as the **designated bridge** between the remote wired subnet and the OWL subnet. 802 frames are bridged from the remote wired subnet to the radio subnet and from the radio subnet to the remote wired subnet through the secondary WDAP. If more than one secondary WDAP is attached to the remote LAN, then the AP with WDAP with the highest bridge priority is "elected" as the designated bridge.

Secondary LAN - an 802 subnet in an 802 LAN which is attached to the OWL network through a secondary WDAP.

Access Point (AP) - a primary, distributed, or secondary access point.

Station - an entity in the 802 LAN which has a unicast 802 address.

OWL Station - a station in an OWL radio network.

Remote Station - a station which is not in an OWL radio network (i.e. a station on the distribution LAN or a secondary LAN).

Node Address - the 48-bit 802 address which uniquely identifies a node in an OWL network.

Node ID - A 16-bit node identifier which can be used to replace the 48-bit node address. In a hierarchical network, each OWL node must obtain a node ID from the super root. The concatenated Net ID and node ID uniquely identify the node within the radio network.

Port Address - the 48-bit 802 address or 16-bit address which uniquely identifies a port in an OWL network. A port address can also be used as the node address.

Port ID - A 16-bit port address, which is used to replace the associated 48-bit port address. In a hierarchical network, each OWL node can obtain a network unique port ID from the super root, for each of its ports. The concatenated Net ID and port ID uniquely identify the port within the radio network.

originator - the node which originates a unicast or multicast transmission.

sink - the target node of a unicast transmission.

conversation - a series of transmissions which are used to forward a frame from an originator to a sink. The frame may be divided into multiple **fragments**.

MDPDU - a MAC-D sub layer protocol data unit.

MRPDU - a MAC-R sub layer protocol data unit.

MQPDU - a MAC-Q sub layer protocol data unit.

MSPDU - a MAC-S sub layer protocol data unit.

MQPDUID - The concatenation of the MQPDUID and 802 source and destination addresses uniquely identifies an MQPDU in an OWL radio network.

inbound - Nodes which are logically closer to the root node of a spanning tree are considered "inbound" from nodes which are further from the root. A **inbound PDU** is any PDU which is traveling toward the root.

outbound - Nodes which are logically further from the root node of a spanning tree are considered "outbound" from nodes which are closer to the root. An **outbound PDU** is any PDU which is traveling away from the root. An OUTBOUND bit in a MAC-R control field is set ON to indicate that the source of a MRPDU is inbound from the destination of the PDU. Note that terminal nodes never set the OUTBOUND bit ON.

OWL Addressing.

Each physical port in an OWL network has a unique 48-bit 802 address. An OWL access point (AP) has an ethernet port and 1 or more radio ports, each with its own unique 802 address. Each radio port requires a unique address, because it is possible for more than one radio port to be active on the same radio channel at the same time. The 802 address for one of the ports is also used as the MAC-R **node address**. A terminal node has a single radio port. The 802 radio port address is also used as the terminal's node address. A 16-bit radio port ID can be used as a 16-bit port address in lieu of a 48-bit 802 port address, for either a terminal or AP. A 16-bit port ID is obtained from an address server in the root node of an OWL LAN. A 16-bit port ID is unique within the domain of an OWL LAN ID. The port address size used on each port (i.e. 2 or 6 bytes) is dependent on the underlying MAC-D type. The address size is assigned when the port is initialized and is fixed for all packets transmitted or received on the port.

The 16-bit port ID for one of the AP ports is also used as the node ID. The node ID for a terminal must be the same as its 16-bit radio port ID. The 16-bit node ID uniquely identifies a node in the OWL radio network. A child's node ID will appear in a pending message list (i.e. in a HELLO packet) when the parent AP is storing messages for the child.

OWL packets carry data and network management information within the OWL radio network. For OWL packets, the DIX ethernet type is hex. 875C or the SNAP SAP is hex. 00C0B2875C. OWL packets contain a MAC-D and a MAC-R header. The destination and source addresses in the MAC-D header are required to uniquely identify the hop destination port and source port, respectively. The destination and source addresses in the MAC-R header identify the end-to-end destination and source nodes, respectively. In addition, inbound unicast MAC-R PDUs which are relayed by an AP, contain the node address of the AP. When an ethernet frame is bridged into the radio network, the detination and source 802 addresses in the ethernet header are copied into the MAC-R destination and source address fields.

An AP bridges packets across subnet boundaries (i.e. radio network to ethernet) and bridges packets from AP ports to internal applications. Since an AP port address may be different than the AP node address, an AP must internally "bridge" packets from an AP port to/from higher layer service access points (SAPs) contained in the AP. The source 802 address for packets generated by a higher layer SAP is always the AP node address. Likewise, the destination 802 address for packets destined for a higher layer SAP in the AP, is the AP node address. For example, the 802 address returned in an (i.e. TCP/IP suite) ARP packet is the AP node address. IP packets destined to the node address can be received on any AP port (i.e. with a different ethernet address) because each port is operating in promiscuous mode. Packets destined to the AP node address are forwarded to the higher layer SAP identified in the DIX ethernet or LLC header. As a second example, assume that an AP has two 2.4 GHz radio ports, each with a port address which is different than the AP node address. A remote RF device can send a packet to the AP node address by using the AP port address (i.e. the MAC-D destination address) to relay the packet. The port address uniquely identifies one of the radio ports. The remote RF device does not need to know that the MAC-R destination exists on the AP that owns the port. Once the packet arrives at the AP, the MAC-R destination address is used to determine that the packet belongs to a local SAP.

MAC-D Sub Layer.

The MAC-D sub layer controls access to the channel and is responsible for providing reliable transmission between any two devices in the radio network. A radio network may include both wired and radio links. The MAC-D sub layer is specific to the physical link type. An 802.3 MAC-D sub layer is used on 802.3 links and a radio MAC-D sub layer is used on radio links.

MAC-D Sub Layer for radio links.

The radio MAC-D sub layer provides "acknowledged connectionless" services to the MAC-R sub layer. A "connection" is not required to transmit an MRPDU; however, each PDU is acknowledged at the MAC-D sub layer and errors are reported to the MAC-R sub layer. For a terminal node, a MAC-D link error provides an indication that the terminal has roamed.

Radio MAC-D Protocol Data Units.

An MDPDU is classified as either a control frame or a data frame. Control frames facilitate network access and error recovery for unicast conversations. Data frames contain an MRPDU. A common header format is used for both control and data frames.

MAC-D header format.

protocol ID
network ID
destination node ID
source node ID
control
reservation

Control frames.

Control frame format.

preamble
SFD (start frame delimiter)
<physical layer header>
MAC-D header
CRC

Note that control frames have a fixed length.

Control frame types.

A control frame is classified as either a request frame or a response frame. A single bit in the type field indicates if a control frame is a request or a response.

Control request frame types.

RTS - an RTS frame is used to reserve the network for a unicast conversation.

ENQ - an ENQ frame is used by an originator to determine the status of a previous fragment transmission. The sink responds by re-transmitting its last ACK or CLEAR. If the sink node does not have ACK state information, it responds to an ENQ by transmitting a REJECT. Note that an ENQ/ACK pair correspond to an RTS/CTS pair with respect to channel access.

ABORT - an ABORT can be used by an originator to abort an active conversation. Note that a conversation can be restarted at any time.

Control response frame types.

CTS - a CTS frame is used to acknowledge an RTS frame and grant access to the network.

ACK - an ACK frame is used to acknowledge the reception of a unicast data frame fragment. The control byte in an ACK frame contains the 1-bit sequence number of the next data frame fragment expected.

CLEAR - a CLEAR frame is used to acknowledge the reception of the last unicast data frame fragment in a conversation. A last-in-chain (LIC) bit distinguishes a CLEAR frame from an ACK frame.

REJECT - a REJECT frame is used by a sink to notify an originator that a unicast conversation has been aborted by the sink or that the sink does not have ACK state information for the originator. The originator must restart the conversation.

Data frames.

Data frames are used to send MAC-R data. The control field in a data frame contains a 1-bit sequence number used to facilitate fragmentation and re-assembly of large unicast frames. All broadcast and multicast transmissions consist of a single DATA frame. Unicast frames may be broken into multiple DATA fragments for transmission. A first-in-chain (FIC) bit is set ON in the first DATA fragment of a frame. A last-in-chain (LIC) bit is set ON in the last DATA fragment of a frame. Note that both FIC and LIC are set ON in single-fragment frames. An EOD (end-of-data) fragment is a data fragment with the LIC bit set ON. Fragmentation and re-assembly at the MAC-D sub layer is transparent to the MAC-R sub layer.

Data frame format.

preamble
SFD
MAC-D header
MRPDU fragment
CRC

Frame transmission.

Multicast frames are sent as a single EOD frame.

Example multicast transmission:

```
EOD ------------------->
```

Example unicast transmission with no errors:

```
RTS --------------------->
   <--------------------- CTS
DATA 0 ------------------>
   <--------------------- ACK 1
DATA 1 ------------------>
   <--------------------- ACK 0
EOD 0 ------------------->
   <--------------------- CLEAR 1
```

If a sink receives an RTS frame and the channel is reserved, then the sink must withhold the CTS frame. The originator must calculate a random back off time and retry later.

Example transmission with errors:

```
RTS ------------------------>
  <------------------------ CTS
DATA 0 ------------------->
  <------------------------ ACK 1
     { ACK lost}
ENQ ------------------------>
  <------------------------ ACK 1
DATA 1 -------------------->
     { DATA lost }
ENQ ------------------------>
  <------------------------ ACK 1
DATA 1 -------------------->
  <------------------------ ACK 0
EOD 0 ---------------------->
  <------------------------ CLEAR 1
     { CLEAR lost }
ENQ ------------------------>
  <------------------------ CLEAR 1
```

Radio Channel Access.

Channel access in an OWL radio network is complicated by the presence of multiple overlapping radio coverage areas and **hidden nodes**. A given first radio transceiver is said to be hidden from a second transceiver, if the second transceiver is not in range of the first, but transceivers exist which are in range of both. In figure 6, the large circles represent the radio coverage area of nodes A, B, C, and D. C, for example, is considered to be hidden from A since it is not in A's coverage area, but a node, B, is in the coverage area of both.

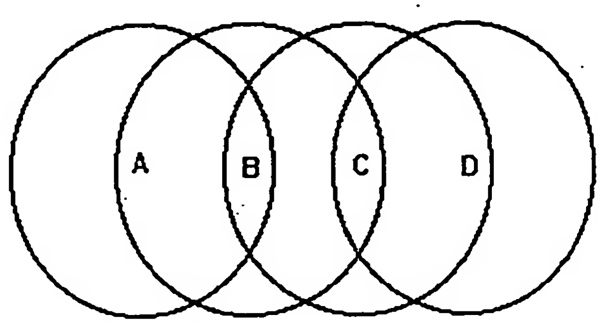

figure 6.

The hidden node problem can severely limit bandwidth utilization in a simple carrier sense radio network if the percentage of hidden nodes is significant. As an example, assume that node A, in figure 6, is transmitting a frame to node B. If, at the same time, C senses the channel it will appear idle, since C can not hear A. If C begins transmitting to D, the transmission from A will collide with the transmission from C at B and will likely be lost. (The transmissions from A and C will not collide at D.)

The OWL MAC-D sub layer uses a listen-before-talk (LBT) collision avoidance protocol to reduce the number of collisions caused by hidden nodes. Nodes reserve the channel for unicast conversations. The reservation in request frames reserves the channel for succeeding data frames. Response frames echo the reservation in the previous corresponding request frame. The reservation in a request frame does not have to span an entire conversation since the reservation can be extended in succeeding data frames. (Shorter reservations reduce dead times when frames are lost.) The reservation in a request frame includes an implicit reservation for the required response (including turnaround time).

The channel reservation technique generally restricts channel access contention to RTS frames. In the absence of lost frames, an LBT algorithm is executed only once per MAC-D conversation. An originator executes the LBT algorithm and transmits an RTS frame if the channel is free. The originator owns the channel for the duration of a conversation as soon as it receives a CTS from the sink. Subsequent DATA fragments can be sent without additional channel access logic. If the channel is not free, a **random back off algorithm**, chooses a back off delay as a function of the LBT slot time and the number of retries. An LBT slot is defined as a function of the best case and worst case busy-sense time. The best case busy sense time is equal to the amount of time from the point at which a node detected the channel idle, before transmitting, until another node can detect the transmission in progress. The worst case busy-sense time is equal to the time required by the originator to sense the channel idle and send an RTS frame plus the time required by a sink to start sending a CTS frame. Figure 7 shows a time line for a unicast conversation between two nodes, A and B. If the originator, A, senses the channel idle at time 0, then the worst-case busy sense time is t_{ws}.

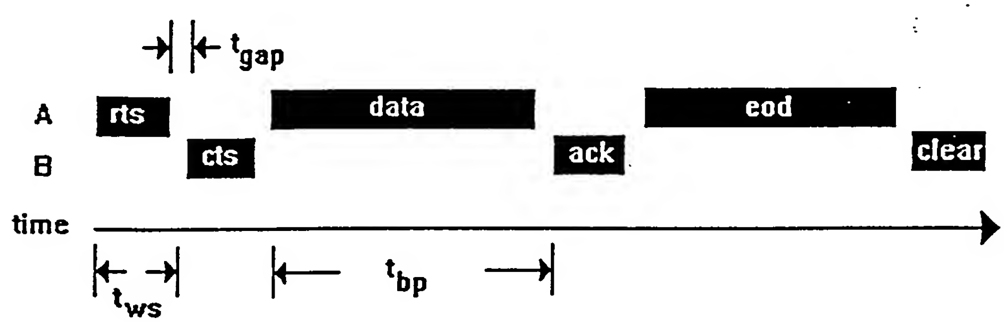

figure 7.

Each node in the network must maintain RESERVE_TIME and RESERVE_NODE channel reservation variables and a local clock. The channel is reserved if the RESERVE_TIME value is greater than the current time. The RESERVE_TIME variable is updated if a reservation is received and 1) the channel is currently not reserved, or 2) the transmitter of a request frame is the RESERVE_NODE node, or 3) the destination of a response frame is the RESERVE_NODE node, or 4) the reservation field in a unicast frame reserves the channel for a time greater than the current RESERVE_TIME period. The RESERVE_TIME is set to 0 whenever a reservation of 0 is observed and the RESERVE_NODE node is the destination of a response frame. The RESERVE_TIME is set to 0 whenever the local node is the target of a unicast transmission from the RESERVE_NODE.

The RESERVE_NODE is set to the concatenated Net ID and node ID of the node which is reserving the channel (i.e. the source node ID and Net ID in a request frame or the destination node ID and Net ID in a response frame) whenever the RESERVE_TIME is updated.

The channel is considered busy if it is sensed busy or if it is reserved. If the channel is reserved, the random delay, chosen by the random back off algorithm, is added to the reserve time. When the delay time expires, the originator repeats the LBT algorithm.

A basic service set (BSS) includes a WMAP and its children. In a frequency hopping network, each BSS is, for the most part, isolated from its neighbors by channel frequency separation, but BSS frequencies may occasionally overlap. Reservations may be missed if BSS frequencies overlap for part of a conversation. If a frequency hop time begins with a synchronization frame, then the synchronization frame can include an indication that the channel is busy.

A sleeping node is any node which has not been actively listening to network traffic. A sleeping node may miss an RTS/CTS sequence. The OWL radio MAC-D protocol uses a busy-pulse technique to support sleeping terminals. CTS and ACK frames provide periodic pulses to indicate that the source node is busy. A sleeping terminal is required to monitor the channel for a busy-pulse period before accessing the channel. If a conversation is in progress, the terminal is guaranteed to hear either the originator or the sink within the busy-pulse period. In figure 7, the busy-pulse period is t_{bp}. The busy-pulse period is well-defined if the maximum fragment and turn-around times are fixed. The combined OWL reservation and busy-pulse protocols provide a channel access solution which is analogous to a busy-tone channel access protocol.

Terminal nodes should limit the total retry time at the MAC-D sub layer, so that roaming can be quickly detected, and a new path in the spanning tree quickly re-established. Relay nodes should lower the number of retries, due to lost frames, when the sink is a terminal node, since the lost frames may be due to roaming. The retry limit should be much higher when both the originator and sink are relay nodes.

802.3 MAC-D Sub Layer.

The 802.3 MAC-D sub layer is used to forward MAC-R PDUs across 802.3 links. All 802.3 MAC-D frames use a common reserved 802 multicast address and LLC SNAP access point identifier in the 802.3 and LLC header, respectively. The OWL MAC-D PDU is contained within the LLC PDU. The 802.3 MAC-D sub layer is used when two (or more) nodes in the OWL network spanning tree are physically connected by an 802.3 link. Note that the same physical link can function both as a distribution LAN and as the physical link associated with a path in the network spanning tree. It is important to understand the following distinction. If a WDAP bridges a frame (i.e. from the radio network) onto a distribution LAN, then the frame is no longer on a branch in the OWL network spanning tree, even if the destination 802 address belongs to a node in the OWL subnet; however, if a WMAP routes an MRPDU to another WMAP then the PDU is forwarded on a branch in the spanning tree, even if the physical link used to forward the PDU also serves as the distribution LAN.

The 802.3 MAC-D PDU fields are shown below. All 802.3 MAC-D transmissions consist of a single data PDU. No control frames are defined. An 802.3 MAC-D sub layer does not fragment MAC-R PDUs.

802.3 MAC-D header format.

protocol ID
network ID
destination node ID
source node ID
control
reservation

802.3 MAC-D data frame format.

802.3 header
LLC header with SNAP access points
MAC-D header
MRPDU
CRC

MAC-R Sub Layer.

The MAC-R sub layer is responsible for correctly routing higher layer PDUs through the OWL subnet. OWL nodes are organized into a network spanning tree and PDUs are routed along branches of the spanning tree. The MAC-R sub layer also provides support for sleeping terminals and distributes network node IDs. The MAC-R sub layer provides unacknowledged connectionless services.

MAC-R Protocol Data Units

MAC-R Header Format

control
destination 802 address
source 802 address
<type specific fields and optional parameters>

MRPDU types.

REGISTRATION - A node sends a REGISTRATION request to the super root to obtain an OWL network node ID, and, optionally, a port ID for each of its physical ports. The registration PDU contains the 802 node address and, optionally, a list of 802 port addresses. The super root records the 802 addresses and returns a node ID, and a port ID for each port address, in a REGISTRATION response PDU. A REGISTRATION request may contain a node alias. The alias is the permanent name of a node in the OWL radio network.

ATTACH - A node sends an ATTACH request to a parent node to attach to the OWL subnet. The ATTACH request is forwarded to the distributed root to establish full connectivity in the OWL subnet. The distributed root returns an ATTACH response packet to acknowledge the ATTACH request. An attach indication (ATTI) bit in the control field of the ATTACH request indicates if the path to the node which generated the ATTACH request has changed. The MAC-R entity in a distributed WDAP (i.e. access root) sets a DISTRIBUTED bit ON in the control field of an ATTACH request before forwarding the request to the super root. The super root records the DISTRIBUTED bit in its routing table and does not bridge frames from the distribution LAN to the attaching node if the DISTRIBUTED bit is ON (i.e. because the distributed WDAP is responsible for bridge frames to the attaching node).

HELLO - Each relay node in a hierarchical OWL radio network periodically broadcasts HELLO response PDUs to advertise its presence. Pending messages for sleeping terminals and broadcast messages can be associated with HELLO PDUs. A node can send a HELLO request PDU to solicit (unscheduled) HELLO response PDUs from attached relay nodes. Each HELLO response PDU contains the 802 address of the super root and a super root sequence number. The super root address and sequence number are used to uniquely identify an occurrence of an OWL network. In addition, each node in the network can learn the 802 address of the super root. Optionally, HELLO response PDUs can contain an encrypted security ID, which uniquely indentifies the OWL domain. HELLO PDUs are ignored if the security ID does not match the OWL domain security ID.

DATA - DATA request MRPDUs are used to transport higher layer data.

R-DATA - DATA response MRPDUs are used to reroute undelivered DATA request MRPDUs after a route has changed.

ALERT - A relay node sends an inbound ALERT request when it is unable to deliver a PDU to a child. The ALERT request is used to determine if the path to the child is still valid and is optionally used to alert the child that it has missed a PDU and should re-attach.

DETACH - A relay node sends a DETACH response node to delete a path to an outbound node.

OWL Network Spanning Tree.

Nodes in an OWL radio network are organized into a network spanning tree. A primary WDAP serves as the (super) root of the spanning tree. PDUs are routed along branches of the spanning tree. Figure 8 shows physical devices and links in an example OWL network. Figure 9 shows the same network organized as a logical network spanning tree.

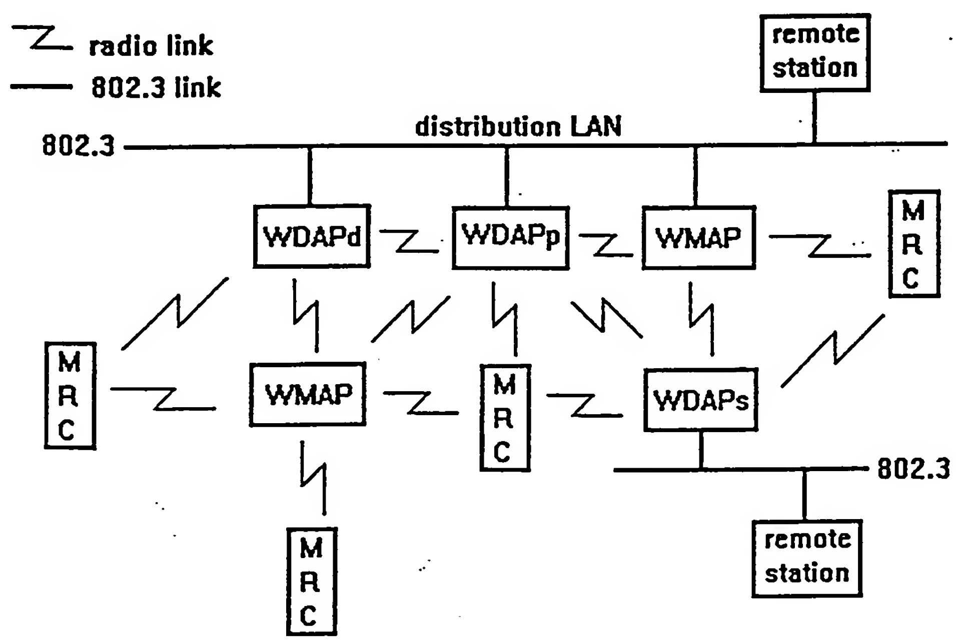

figure 8.

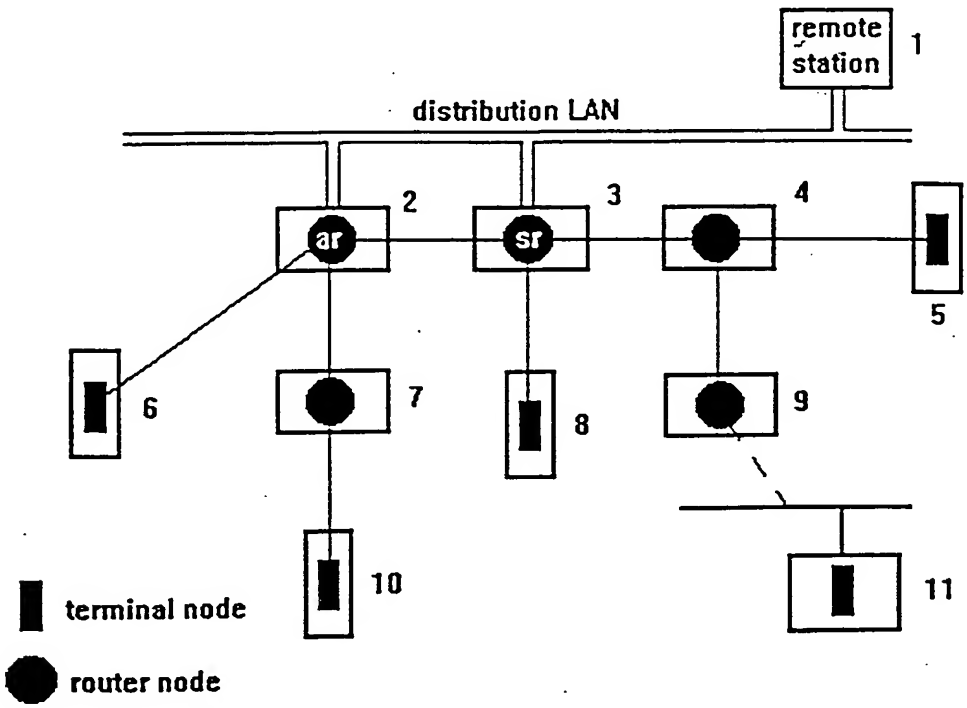

figure 9.

The spanning tree eliminates loops in the physical topology. The node labeled "sr", in figure 9, is the super root and the node labeled "ar" is an access root. The parallel lines represent the distribution LAN, which is not part of the spanning tree. The super root and access root both have access to the distribution LAN; the WMAP labeled 4 can not directly access the distribution LAN. WMAP 4 forwards PDUs destined for the distribution LAN through the super root (i.e. with an 802.3 MAC-D sub layer). The remote station, labeled 1, on the distribution LAN is not part of the network spanning tree; however, the secondary 802 LAN and the remote station, labeled 11, can be viewed as part of the spanning tree (as indicated by the dotted edge).

Building the Spanning Tree.

Nodes in the radio network are generally categorized as attached or unattached (i.e. to the network spanning tree). Initially, only the super root is attached. A single WDAP can be designated to contain the root node, or multiple root candidates can negotiate to determine which node assumes the super root status. The root and other attached relay nodes broadcast HELLO response PDUs at calculated intervals. The HELLO response PDUs enable unattached nodes to learn the optimum path to the super root before attaching to the network. The HELLO response PDUs include: 1) the source node ID and 802 address; 2) a broadcast destination node ID and 802 address; 3) the "cost" to the super root; 4) a "seed" value used to calculate the time of the next HELLO response PDU; 5) a hello displacement time; 6) the priority of the super root node (or root candidate); 7) the 802 address of the super root (or root candidate); 8) a super root sequence number, used to distinguish between multiple occurrences of the network spanning tree with the same super root; and 9) an optional security ID, which uniquely indentifes the OWL domain.

The HELLO "cost" field indicates the total "distance" to the super root, and is equal to the sum of the costs of each hop on the path to the root. (Note that the super root broadcasts HELLO PDUs with the cost field set to zero.) The incremental cost of the hop between a node and its parent is primarily a function of the physical link type (i.e. ethernet or radio). The cost component is intended to bias path selection toward high-speed (i.e. wired) connections. On radio links, spanning tree attachment is biased toward the link

with the best signal strength. Signal strength is not a factor in the cumulative path cost. The HELLO "displacement" field specifies the displacement of the actual hello time from the calculated hello time or indicates that the hello time was unscheduled. A well-known randomization algorithm is used to calculate the next hello time. The HELLO "seed" field is used as a seed for the calculation. The "root 802 address" and "root sequence" fields are used to define a single instance of the radio network. Attached nodes must forget their node ID and return to the unattached state whenever a HELLO response PDU is received with a new root 802 address or root sequence number. HELLO response packets can contain other optional parameters such as a load indication, a distributed clock or a pending message list. A load indication parameter can be used to balance the number of terminal nodes between access points.

Nodes without a parent in the spanning tree are in an **unattached** state. In the unattached state, a node learns which attached relay node is closest to the super root by listening to HELLO response PDUs. (If no HELLO response PDUs are received, the node can wait (i.e. sleep) and retry later.) After the learning period expires an unattached node sends an ATTACH request packet to the attached relay node with the lowest cost to the super root. The ATTACH request contains an ATTACH ID, which is a sequence number that is incremented whenever an ATTACH request is generated. (Nodes without a node ID must first send a REGISTRATION request packet to the root to obtain an OWL node ID.) The attached relay node forwards the ATTACH request packet to the super root. <u>The end-to-end ATTACH request functions as a discovery packet and enables relay nodes along the path to the super root to quickly learn the path to the source node</u>. The super root returns the request as an end-to-end ATTACH response PDU. <u>The node which originates an ATTACH request is responsible for retrying the request until a matching response is received, to insure that it is fully attached.</u> When the unattached node receives the ATTACH response PDU it goes to an **attached** state and sets internal **root port** and **parent** variables. The root port is the physical port on which the response PDU arrived and the parent variable contains the node ID and 802 address of the parent node. <u>A child node will only accept outbound unicast MRPDUs from its parent.</u> If the newly attached node is a relay node, it calculates its cost to the super root, by adding its root port link cost to the HELLO cost of its new parent, and begins to broadcast HELLO response PDUs.

ATTACH requests are always forwarded to the super root. Inbound ATTACH requests establish a new path from the super root to the source node. The super root will convert an ATTACH request to either an ATTACH response or to a DETACH response (i.e. if an old path exists). In general, if an ATTACH response is at an AP, and a AP exists which is on the old path to the attaching node, but is not on the new path, then the AP must convert the ATTACH PDU to an outbound DETACH request PDU and forward it on the old path. When the "old parent" AP receives the DETACH request, it will read filter and forward sequence numbers from the request and will enter its filter and forward sequence numbers for the source node into the request. The old parent will then delete its routing table entry for the node which originated the ATTACH request and will return the DETACH request as an inbound DETACH response. The AP which originated the DETACH request will convert the DETACH response to an outbound ATTACH response and will forward it to the attaching node. The old parent AP can (optionally) re-route any undelivered PDUs, destined to the attaching node, as MAC-R R-DATA PDUs. An R-DATA PDU is routed inbound until the new outbound branch is reached. If the new branch has been used, the R-DATA PDU is discarded (i.e. to avoid out-of-sequence delivery); otherwise, the PDU is routed outbound along the new path.

Unattached terminal nodes can optionally broadcast a global HELLO request PDU with a multicast relay node ID and broadcast 802 destination address to solicit unscheduled HELLO response PDUs from attached relay nodes. The net effect is that the unattached state can (optionally) be shortened. (Note that only attached relay nodes respond to request PDUs.) The HELLO request facility is intended for unattached terminals with higher-layer transmit requests in progress.

Each attached node must transmit an ATTACH request PDU at least once per ATTACH_TIMEOUT time period to maintain its path in the radio network. An attached node must also transmit an ATTACH request PDU to its parent whenever it misses MAX_HELLO_LOST consecutive scheduled HELLO response PDUs from its parent and whenever it receives an alert. An alert can occur in an alert PDU or in an

optional alert list in a HELLO response PDU. If a relay node is unable to deliver a PDU to a child node, then the relay node adds the node ID of a child node to its alert node list and, optionally, generates an alert PDU which is sent down all branches of the spanning tree.

Each node (except the super root) should maintain an **in-range list** which contains the node ID and 802 address of potential alternate parent nodes. If a child loses its parent (i.e. due to a MAC-D link error) or detects a better path, then the child can change its path in the spanning tree by selecting the best candidate from the in-range list and attaching to the new parent. Relay nodes must avoid sporadic path changes. If a child loses its parent and the in-range list is empty, it must remain in a quiet learning state until a potential parent is discovered.

MAC-R Routing.

All PDUs are routed along branches of the spanning tree. Relay nodes "learn" the path to outbound nodes by monitoring inbound traffic (i.e. traffic directed toward the root). Whenever a relay node receives an inbound REGISTRATION or ATTACH request PDU from an outbound node, it creates or updates an outbound entry for the source node in its **routing table.** The entry includes the source node's 802 address and the port address of the port which sent the PDU (i.e. the hop address). When a relay node receives a PDU from an inbound node the PDU is forwarded to the outbound hop address which is specified in the routing entry for the 802 destination. The PDU is discarded if a routing entry does not exist. (Note that an AP may also maintain a filtering database which contains entries for inbound nodes and nodes on either the distribution LAN or a secondary LAN. The filtering database is used to facilitate bridge layer flooding.)

As an example, the routing table for relay node 4, in figure 9, is shown in figure 10 below. The **destination** field contains the 802 address of a node in the sub tree rooted at 4. The **first hop** field contains the MAC-D (i.e. port) address of the first hop on the path to the destination. (The node labels from figure 9 are used in lieu of node and port addresses, in this example.) The **child** field indicates if the destination is a child. The **attach ID** field is used to associate ATTACH and DETACH requests and responses. The **port** field specifies the physical port used to communicate with the first hop. The **type** field can be RELAY or TERMINAL. The **status** field is used to mark each entry as BOUND or UNBOUND. UNBOUND entries can be used to route outbound REGISTRATION response PDUs. An UNBOUND entry becomes BOUND when the AP receives an ATTACH response PDU with a matching Attach ID. The super root must also mark each entry which specifies a path through a distributed WDAP as DISTRIBUTED. The **age** field indicates the last time the destination was active and is used to "age" away old table entries. Assume that relay 4 has received an ATTACH request from node 11 through relay 9. Relay 4 adds an entry for destination 11 with the first hop set to 9, the age set to 0.

Destination	Type	Child	First Hop	Attach Time	Attach ID	Port	Status	Age
11	TERMINAL	No	9	1223	4	1	ATTACHED	0
5	TERMINAL	Yes	5	802	2	1	ATTACHED	2
9	RELAY	Yes	9	907	5	1	ATTACHED	1

figure 10.

PDUs from outbound nodes are forwarded to the next inbound node (i.e. the parent) in the branch of the spanning tree. <u>No explicit routing is required for inbound traffic because the route is defined by the structure of the spanning tree</u>. A PDU travels inbound until a node is reached which has an entry in its routing table for the destination 802 address. The PDU is then explicitly routed outbound until it reaches

its destination. Thus, communications between any two nodes is accomplished by routing all traffic through the nearest common ancestor of both the source and destination node. If a PDU reaches a primary or distributed WDAP and an entry for the 802 destination does not exist in the routing table of the WDAP, then the PDU can not be routed outbound (i.e. a common ancestor does not exist). In this case, the WDAP can "bridge" the PDU, as an 802 frame, onto the distribution LAN. Note that a PDU which is bridged onto the distribution LAN by a distributed WDAP, will be bridged back into the OWL subnet (i.e. by another WDAP) if the 802 destination is in the OWL subnet.

As an example, in figure 9, if a PDU is sent from terminal 10 to terminal 5 it will be routed as follows: Terminal 10 will send the PDU to its parent, WMAP 7. Since WMAP 7 does not have an entry in its routing table for terminal 5, it will forward the PDU inbound to its parent, WDAP 2. The MAC-R entity in WDAP 2 does not have an entry in its routing table, so it will forward the PDU to its bridging entity and the PDU will be bridged onto the distribution LAN as an 802 frame. The bridging entity in WDAP 3, the super root, , will forward the frame to its MAC-R entity because it has an entry in its forwarding data base, which specifies the radio network as the subnet for terminal 5. The MAC-R entity in WDAP 3 has an entry in its routing table for terminal 5 and will forward the PDU to the first outbound hop, WDAP 4, over the wired link (i.e. with an 802.3 MAC-D sub layer). WDAP 4 will then deliver the PDU to terminal 5.

As a second example, if remote station 11, in figure 9, sends a PDU to remote station 1 it will be routed as follows: The bridging entity in the secondary WDAP, 9, will determine that station 1 is not on its local 802.3 subnet (i.e. by querying its forwarding database) and will bridge the PDU into the radio network (i.e. by passing the frame to its MAC-R entity). The MAC-R entity in WDAP 9 will forward the PDU inbound to WMAP 4, since it does not have an entry for station 1 in its routing table. WMAP 4 will forward the PDU to WDAP 3. The MAC-R entity, in WDAP 3, does not have an entry for station 1 and will pass the PDU to its bridging entity. The bridging entity will forward the PDU onto the distribution LAN as an 802 frame addressed to station 1.

Dynamic routing changes and PDU retransmission.

Paths in the spanning tree change often as terminals roam. PDU transmission errors due to roaming fit into one of two possible cases: 1) a terminal node is unable to deliver a PDU to its parent AP, or 2) an AP is unable to deliver a PDU to a child terminal.

In the first case, the terminal can simply select a new parent and re-attach to the network by sending an ATTACH request. An attach indication is generated whenever the path to a terminal node changes. The MAC-R entity in a relay node updates its routing table entry for an outbound source node if an inbound ATTACH (or REGISTRATION) request PDU is received from the node and the hop source is not the same as the first hop in the table entry for the node. The first hop field, in the routing table entry, is overlaid by the hop source of the PDU and outbound PDUs are now routed along the new path. (Note that an old disconnected path fragment may still exist in the spanning tree after a new path has been established.) ATTACH requests are always forwarded to the super root.

An ATTACH PDU may be converted to an outbound DETACH request to deleted an old path fragment. The ATTACH ID in a DETACH request is the same as the ID in the associated ATTACH PDU and the destination is the 802 address of the attaching node. If a relay node on the old path has a routing table entry for the destination, with an ATTACH ID that matches the ID in the DETACH request, then the relay node will forward the DETACH request outbound and will delete the entry. The DETACH request is forwarded outbound until it reaches the relay node which was the old parent of the attaching node. It is then converted into a DETACH response and is forwarded inbound until it reaches an AP which is on the new path to the attaching node. The AP on the new path will convert the DETACH response to an ATTACH response and route the ATTACH response outbound to the attaching node. Note that if the ATTACH response is lost, then a DETACH PDU will not be generated on the next ATTACH retry (i.e. because the old path was deleted).

A relay node may not be able to deliver a DATA PDU to a child, for several reasons: 1) the child may be asleep; 2) the channel may be reserved in the child's coverage area; 3) the PDU may be lost due to excessive errors; or 4) the child may have selected a new parent (i.e. due to roaming). It is assumed that most undelivered PDUs are lost because child nodes roam. If a parent relay node can not deliver a PDU to a child node, then (if the routing table entry for the child has not been updated) the parent node will add an alert record for the child node to its internal alert list and send an ALERT request to the super root. The ALERT PDU contains the ATTACH ID from the routing table entry for the child node. When a relay node, on the path to the super root, receives an inbound ALERT request it determines a) if the alert ATTACH ID matches the ATTACH ID in its routing table and b) if the hop source in the ALERT request is the same as the first hop field in the routing table entry for the alert destination. If both conditions are satisfied then the relay node will 1) optionally add the associated alert record to its internal alert list, 2) forward the ALERT request to the next hop on the path to the super root, and 3) optionally forward the ALERT request down each of its outbound branches, other than the one on which it arrived. If either condition is not satisfied then the relay node will, instead, send an outbound ALERT response on the path on which the ALERT request arrived. The ATTACH ID in the ALERT response is the same as the ID in the ALERT request and the destination is the 802 address of the lost child. If a relay node on the old path has a routing table entry for the destination, with an ATTACH ID that matches the ID in the ALERT response, then the relay node will forward the ALERT response outbound and will delete the entry. The ALERT response is forwarded until it reaches the relay node which was the old parent of the lost child.

An ALERT request may reach the super root before the associated child node re-attaches. In this case, the ALERT request is simply discarded.

Outbound ALERT requests are used to quickly notify a lost child that it should re-attach to the network. If a relay node receives an outbound ALERT (i.e. from its parent) request, it first checks to see if it has a routing table entry for the lost child with a "newer" ATTACH ID. If it does, then the ALERT request is simply discarded. Otherwise, a relay node which receives an outbound ALERT request will forward the ALERT request to each child node which is a relay node and will multicast the ALERT request (i.e. with a multicast MAC-D destination address) once on each of its radio ports. Each relay node adds the ALERT ID in the request to its internal alert list

Records in a relay node's internal alert list in each relay node are copied into HELLO response PDUs for MAX_HELLO_LOST + 1 scheduled hello times to notify nodes to re-attach, where MAX_HELLO_LOST is the maximum number of HELLO PDUs that can be missed by a child before the child re-attaches. An alert record contains a target node ID, a source node ID, and an ALERT ID (which equates to an ATTACH ID). The concatenated source node ID and ALERT ID are used to uniquely identify each alert occurence. A target node can ignore any any duplicate alert record which is received within MAX_HELLO_LOST+5 HELLO periods.

Note that old path fragments are simply aged away if no outbound PDUs are sent along the path fragment.

A terminal node must set the attach indication (ATTI) bit ON in the MAC-R header of an ATTACH request when it first attaches to a new parent. The ATTI bit indicates that the path to the ATTACH request source node has changed. The MAC-R entity in the AP which was the previous parent of the source node posts an attach indication error to the MAC-Q sub layer when it receives the ATTACH request PDU with the ATTI bit set ON. The MAC-R sub layer in a terminal node posts an attach indication error to the MAC-Q sub layer when it receives the associated ATTACH response with the ATTI bit set ON. An attach indication is a positive indication that a node has just attached to the network and can be used to trigger an immediate (re)transmission. The attach indication includes the 802 source address and receive sequence number for the source node of the ATTACH request. If the MAC-Q entity is holding any undelivered DATA PDUS for the node, it can respond by re-transmitting the undelivered PDUs as R-DATA PDUs. The R-DATA PDUs will be discarded if they are duplicates or arrive out-of-sequence. The R-DATA PDUs are automatically routed along the new path.

The MAC-R layer in a terminal node is responsible for retrying a DATA PDU transmission, if the MAC-D layer is unable to deliver the DATA PDU to its parent. The MAC-D layer indicates the success or failure of a transmission. Occasionally, the MAC-D entity will not be able to positively determine success or failure (i.e. if CLEAR frames are missed in a MAC-D conversation). If the MAC-D layer indicates positive failure, then the MAC-R layer can choose a (possibly new) parent, re-attach, and retransmit the DATA PDU; otherwise, the MAC-R layer must discard the PDU. The MAC-Q may retransmit the DATA PDU as an R-DATA PDU when an attach indication is received (i.e. when an ATTACH response is received with the ATTI bit set ON).

Registration.

A node is initially in an unregistered state and returns to the unregistered state under certain error conditions. Each unregistered node in the network must send a REGISTRATION request to the super root before it attaches. The REGISTRATION request is used to obtain a network node ID and is used to validate access to the network. The REGISTRATION request is returned by the super root as a REGISTRATION response. The node which originated the request is responsible for retrying the request until a matching response is received.

Registration logic is similar to attach logic with some key differeneces. REGISTRATION requests can only be sent to the super root when no other inbound PDU for the source node exists in the network. No other PDU types may be sent in the unregistered state. A node goes to the registered state when a matching registration response is received from its parent.

A node's registration is valid as long as it is actively attached to the network. A node returns to the unregistered state if it does not receive an ATTACH response within a MAX_ADDRESS_LIFETIME time period or if it detects that the super root has changed.

Broadcast routing.

PDUs with broadcast (or multicast) 802 destination addresses are (optionally) routed along all branches of the network spanning tree. Broadcast messages are associated with HELLO response PDUs on radio links. A broadcast parameter in a HELLO response PDU indicates that terminals should stay awake for broadcast messages which will immediately follow the HELLO PDU. A secondary WDAP forwards broadcast messages onto its attached wired subnets. If a broadcast message orginates on the distribution LAN, then each primary or distributed WDAP is responsible for bridging it to the OWL sub tree for which it is the access root. Broadcast messages which originate within an OWL subnet are forwarded on each branch of the network spanning tree, except the branch on which the message arrived. The access root of the sub tree in which the broadcast message originated is responsible for bridging the message onto the distribution LAN. The message is (optionally) bridged back into the radio network by each other access root.

As an option, broadcast (or multicast) messages which originate in the radio network are only forwarded to the distribution LAN. A broadcast message which originates in the radio network is forwarded inbound until it is bridged onto the distribution LAN by a primary or distributed WDAP. The message is not relayed back into the radio network by each other primary or distributed WDAP (i.e. as above). To facilitate this option, each distributed WDAP must keep a table of inbound entries. An **inbound entry** is defined relative to a WDAP and is any OWL node which is not in the subtree rooted at the WDAP. A distributed WDAP adds an entry to its inbound table when it receives a DIST_ATTACH packet. A DIST_ATTACH PDU is generated by the access root, which is responsible for bridging to a node, when the node attaches to the network. It is used to notify transparent bridges and other distributed WDAPs that the node has roamed. A distributed WDAP will not bridge a broadcast packet from the distribution LAN into the radio network if the source address belongs to a node in its inbound table. The primary WDAP will not bridge a broadcast packet from the distribution LAN into the radio network if the source node belongs to a node in its route table.

Sleeping Terminal Support.

The MAC-R sub layer provides several facilities to support sleeping terminals. A sleeping node initially "synchronizes" on a HELLO response PDU from its parent. The node can calculate the time of the next expected HELLO response PDU from its parent and can power-down with an active timer interrupt set to wake it just before the next HELLO response PDU is transmitted. The MAC-R entity in a parent node can store a message for a sleeping node until the node "requests" the message by notifying its parent that it is awake. A terminal learns that it must request unsolicited saved messages by examining a pending message list in the HELLO response PDU. This implementation enables sleeping terminals to receive unsolicited messages and relaxes the timing constraints for transaction oriented messages. ATTACH and DATA request PDUs can contain several MAC-R parameters which are used to enable pending messages. A "delivery service type" parameter, indicates that a terminal (i.e. which sent the request) is sleeping. An "awake time window" parameter is used to specify an awake time period. An "awake time offset" parameter is used to specify the start of the awake time window. (The awake time window is effective immediately if an awake time offset is not specified.) An "auto awake" delivery service type can be used to implicitly set an awake time window each time the parent node receives a message from the sleeping terminal. A "maximum stored message count" field specifies the maximum number of HELLO times that a message should be stored in the parent relay node. The MAC-R entity in a parent node will store pending messages until 1) the message is delivered, or 2) "maximum stored message count" hello times have expired.

Broadcast messages are associated with HELLO PDUs so that sleeping terminals will be awake when the broadcast message is transmitted.

WDAP bridging.

A WDAP maintains a forwarding data base with an entry for each known network node. Each entry contains an 802 destination address and an associated subnet identifier. When a PDU arrives at the bridging entity in a WDAP, the forwarding database is searched to determine the subnet of the 802 destination. If the destination is found and the destination is on another subnet (i.e. other than the one on which the PDU arrived) then the PDU is bridged to the subnet of the destination. If the destination is not found, then the action taken by the bridging entity is dependent on the configuration of the WDAP. 1) The PDU can be forwarded to every subnet except the subnet on which it arrived (i.e. flooding), or 2) the PDU can be discarded.

Typically a primary or distributed WDAP is configured to only forward unicast frames from the distribution LAN to the OWL subnet if an entry exists in its MAC-R routing table for the 802 destination. This implies that the MAC-R entity must notify the bridging entity that a destination exists in the radio subnet, when a MAC-R routing table entry is created, so that the bridging entity can update its forwarding database. Likewise, the bridging entity must be notified when a routing table entry is deleted. The forwarding database in a distributed WDAP contains entries for each node in its access spanning tree. The forwarding database in the primary WDAP contains entries for all nodes in the OWL subnet which are not in an access sub tree rooted by a distributed WDAP.

Flooding options.

The user can set unicast and/or multicast flooding options for the distribution system and for each secondary LAN. Distribution flooding options are configured on the primary WDAP and are distributed in REGISTRATION response packets. Therefore, distribution flooding options should be configured for each AP with a non-zero root priority. Multicast and/or unicast flooding should be enabled for each secondary LAN which requires multicast or unicast flooding, respectively. Therefore, multicast and

unicast flooding should be configured on each AP which can be the designated secondary WDAP for a secondary LAN.

If distribution unicast flooding is set to level 1 (the default) then unicast frames which originate on the distribution LAN are discarded if the destination is unknown. Unicast frames which originate in the radio network are forwarded inbound, until the packet arrives at an AP which a) has a route entry for the destination or b) is a primary or distributed WDAP. A primary or distributed WDAP will relay an inbound unicast frame onto the distribution LAN, if the destination is unknown.

If distribution unicast flooding is set to level 2, then MAC-R will flood unicast frames, for which the destination is unkown, to the distribution LAN and to each secondary LAN which has unicast flooding enabled. For example, a unicast frame which originates in the radio network is forwarded to the distribution LAN and to each secondary LAN which has unicast flooding enabled.

If distribution multicast flooding is set to level 1, then multicast frames which originate on the distribution LAN are forwarded to secondary LANs which have multicast flooding enabled. Multicast frames which originate in the radio network or on a secondary LAN are forwarded to the distribution LAN.

If distribution multicast flooding is set to level 2 (the default) then multicast frames which originate on the distribution LAN are flooded throughout the OWL network. Multicast frames which originate in the radio network or on a secondary LAN are forwarded to the distribution LAN.

If distribution multicast flooding is set to level 3 then all multicast frames are flooded throughout the OWL network. For example, if a multicast frame originates in the radio network then it will be forwarded to the distribution LAN and each AP in the OWL network. An AP will broadcast the message on each of its radio ports and will relay the message to any attached secondary LAN.

Frame filters.

A frame filter can be used to filter multicast frames and/or specified frame types. The user can enter a unicast and/or multicast "frame filter expression" for each AP port which is attached to a remote subnet. Each received frame and an expression pointer are passed to **fltr_is_enabled**. TRUE is returned if the frame is enabled; otherwise, FALSE is returned. If the frame is enabled, it will be bridged into the radio network; otherwise, it will be discarded.

Bridging to a remote subnet.

A remote 802 subnet is bridged to the OWL radio network through a secondary WDAP. The secondary WDAP is responsible for attaching the remote subnet and its remote stations to the radio network. Remote stations can be attached in two ways: 1) The AP which is the designated bridge for the secondary subnet (i.e. the secondary WDAP) can include a remote attach list in its ATTACH request packets. 2) The path to a remote station is automatically established whenever the station sends a DATA packet inbound. A remote attach request is generated whenever an AP receives an inbound DATA packet from a remote station and the source is not in the AP's MAC-R routing table. Data can be piggybacked on a remote attach request. A remote attach request is converted to a remote detach request, if the path to remote node changes.

The bridging layer in a secondary WDAP maintains a list of remote stations which exist on the remote subnet (i.e. a forwarding database). The WDAP port attached to the remote subnet always operates in promiscuous mode. If a frame is received and the destination is not in the port's station list, then the frame is forwarded to the MAC-R layer. The MAC-R layer will either forward the frame outbound, if an entry for the destination exists in its routing table, or will forward the frame inbound to its parent.

A secondary WDAP negotiation protocol is used to select a single designated bridge, if more than one secondary WDAP is connected to a wired secondary LAN. The designated bridge is solely responsible for bridging between its secondary LAN and the radio network.

Optimization considerations.

If a primary or distributed WDAP has two subnets - a distribution LAN and the OWL subnet - and the WDAP is configured to allow flooding onto the distribution LAN and to not allow flooding onto the OWL subnet, then each entry in its forwarding database corresponds to an entry in its MAC-R routing table. All frames which are passed to the bridging entity from the MAC-R entity (i.e. from the OWL subnet), are forwarded to the distribution LAN. Frames will only be forwarded from the distribution LAN to the OWL subnet if an entry exists in the MAC-R routing table. For any configuration, entries in the forwarding database which are associated with the OWL subnet correspond to entries in the MAC-R routing table. A shared forwarding database/MAC-R routing table data structure could be used to optimize the learning process required for bridging and to avoid two lookups (i.e. a forwarding database lookup and a MAC-R routing table lookup) each time a PDU is forwarded from the distribution LAN into the OWL subnet.

MAC-Q Sub layer.

The (optional) MAC-Q can be viewed as an end-to-end reliability layer between entry points to the radio network. The MAC-Q sub layer is responsible for delivering received PDUs to the next higher layer in the order in which the PDUs entered the radio network. The MAC-Q sub layer also retransmits lost MQPDUs, and filters any resulting duplicate or out-of-sequence MQPDUs. The MAC-Q sub layer is intended to significantly reduce the number of lost PDUs due to "roaming" terminals, without introducing duplicate or out-of-sequence PDUs. It does not guarantee that PDUs will never be lost. MAC-Q entities exist at entry points to the radio network. The MAC-Q entity in an AP provides a proxy MAC-Q layer for nodes in the OWL network which are not in the radio network.

MQPDUs contain a MQPDUID in the MQPDU header. The concatenation of the MQPDUID and 802 source and destination addresses uniquely identifies an MQPDU in an OWL radio network. The MQPDUID is generated by the MAC-Q entity in a WDAP or terminal when a frame first enters the OWL radio network.

A primary or distributed WDAP maintains an "MQPDU table" with entries for each outbound node. Each entry contains the 802 address of an outbound node and an associated forward MQPDUID and filter MQPDUID. Forward MQPDUIDs are generated to uniquely identify an MQPDU for its lifetime in the OWL network. Filter MQPDUIDs are used to detect duplicate and out-of-sequence PDUs. Before a primary or distributed WDAP forwards an 802 frame from a wired backbone into the OWL radio network, it increments the forward MQPDUID, associated with the destination 802 address, and enters the it into the MQPDU header. The MQPDU is then passed to the MAC-R sub layer for transmission. Note that this approach assumes that remote stations do not move quickly from subnet to subnet. <u>If a node is physically attached to two subnets, then a unique 802 address should be used for each subnet.</u>

Terminal nodes maintain an MQPDU table with an entry for each active remote MAC-Q network entry point. Each entry contains a filter MQPDUID, a subnet identifier, and an 802 address. Subnet 0 is always the radio network and subnet 1 is the distribution LAN. Other subnet identifiers can be assigned to a secondary WDAP. The 802 address is blank for subnets 1 and higher. Note that there can be multiple entries for subnet 0, but only 1 entry for each other subnet. A terminal also maintains a single forward MQPDUID variable and stores up to one MQPDU for possible retransmission. The value of the forward variable is incremented and entered into the MQPDU header whenever a terminal prepares a new PDU for transmission. The terminal MAC-Q entity retransmits an MQPDU whenever the MAC-R layer returns a transmit error (until a maximum retry count is exceeded).

The filter MQPDUID, in an MQPDU table, is the ID of the last MQPDU received from the associated 802 address. Duplicate MQPDUs are discarded. An MQPDU is accepted by a sink if 1) a retry bit in the MAC-Q header is set OFF or if 2) the MQPDUID in the PDU is not in a "duplicate range" defined by the filter MQPDUID in the table. If an MQPDU table filter entry does not exist for an 802 source address, then data PDUs from the source should be discarded if the retry bit is set ON. The entries in the MQPDU table must be aged so that a filter MQPDUID (and stored MQPDU) is never older than the "roll over" time of an MQPDUID.

An entry in an MQPDU table in a distributed WDAP may be transferred to another primary or distributed WDAP if a terminal "roams". If a terminal moves and its new path to the super root is through another WDAP, then the forward and filter MQPDUIDs for the terminal must be transferred from the old WDAP to the new WDAP. The super root obtains the information (if it exists) from the old WDAP and forwards it to the new WDAP. Note that the new WDAP can accept MQPDUs with the retry bit set OFF while waiting for an MQPDU table entry to be transferred.

Ideally, each MAC-Q entity in the radio network should be notified when the terminal node associated with an entry in its forward list has roamed and re-attached. If a MAC-Q entity holds an undelivered PDU, destined for the re-attached terminal, then the PDU can be retransmitted along the new path to the terminal. A more practical approach would be to notify each MAC-Q entity which has recently transmitted a PDU to the terminal. If it is assumed that most traffic is not contained in the radio network, but rather is directed to or from the distribution LAN, then it may be practical to simply notify the MAC-Q entities in primary or distributed WDAPs on the old path to the terminal.

MAC-S Sub Layer.

The (optional) MAC-S sub layer provides data compression and security services.

Network management tools can be used to create security associations between any two stations in an 802 LAN which contains an owl subnet. MAC-S entities exist in WDAP's. A MAC-S entity can encipher a frame when it enters the radio network if a security association exists between the source and destination stations at the entry WDAP. A MAC-S entity, in an exit WDAP, can correctly decipher a frame as it exits the radio network if it contains a corresponding security association. Network management access to a MAC-S entity in a distributed WDAP is always through a primary WDAP. The primary WDAP (i.e. the super root) "knows" the path to all outbound nodes. A MAC-S entity in a primary or secondary WDAP provides a "proxy" MAC-S layer for security associations involving remote stations on wired subnets.

A global security association can be used to consistently encipher and decipher each frame as it, respectively, enters and exits the radio network. Global association must be enabled at the MAC-S entity in each primary, secondary, and terminal node in the OWL subnet.

Simple compression (i.e. independent of any security encryption) is enabled by a single compression bit in the MAC-S header.